PERFORMANCE MODEL BRIEFING

THE ROLE OF SLA MODELS IN BUSINESS MANAGEMENT

Nick Gill, Steve Hodson, 31 May 2005

TABLE OF CONTENTS

OUTSOURCING – AN SLA PERSPECTIVE

There is an increasing trend among companies and public bodies towards the outsourcing of major parts of their infrastructure to specialist service providers. The justification for this is based on many factors such as the ability to maintain and develop skills; the ability to generate economies of scale; and the ability to offer greater career options to the staff being outsourced. In practice there is generally only one over-riding driver and that is cost reduction.

The Service Level Agreement (SLA) is a familiar term in the industry but often misunderstood and incorrectly applied. This paper provides background on SLAs and their role in business management where a service provider has a major responsibility to the business and indeed in an outsourcing role.

The role of the model is to provide a quantitative and factual perspective on the service from provider to client. Outsourcing has become very popular in public and private sectors alike. There are regular accounts of how the expectations of parties have not been met. Whether embarking on an outsourcing project or managing an established one, the model approach enables both parties to share a systematic and structured view of the relationship and analyse the effects of the SLA. By doing this, they can reflect on the effects in terms of incentives to provide the desired service, financial penalties and the way in which the inherent design of the solution is suited to the client's service aspirations.

Intercai has gained substantial experience in SLA education, design and the provision of models to support them. This has been initially been between telecommunications operators and large client organisations. We have applied our skills in an ICT environment and have supported outsourcing providers in major contracts in public and private sectors.

ROLE OF SERVICE LEVELS IN PROCUREMENT

MANAGEMENT AND MISMANAGEMENT

Having made the decision to outsource, the concern arises that the service provider may not deliver the degree of service that the company has enjoyed hitherto, and so a series of performance standards are introduced that are intended to keep the service provider up to the mark during the tenure of the contract. These are intended to manipulate the behaviour of the service provider to maximise the advantage gained by the user of the service from outsourcing it. These standards are normally backed up with contract penalties in the event of failure to meet them. It is common practice that the extent of the penalties could have major impact on the business of the service provider. It is also, strangely, common practice that service providers enter into these contracts with little or no understanding of the risk to the business that the service levels represent and no allowance in their business plan for some level of penalty payment. It is also the case that adverse penalties can act as a disincentive to the supplier.

The original intent of the service levels, laudable as it is, often gets lost in a maze of complexity that micro-manages the service provider to such an extent that there are major and unforeseen consequences – often having the direct opposite incentive to that intended. This micro management also limits the extent to which the service provider can derive the benefits described above with consequent damage both to the business of the service provider and the customer.

SLAS AND RISK MANAGEMENT

It is a general expectation, encouraged by the industry and the trade bodies that have formed around the issue of service levels, that the Service Level Agreement (SLA) represents the service level that will be delivered. This, unfortunately, is a myth. If an SLA sets a threshold at the same level as the average performance of a service then the supplier will, by definition, fail to meet it half of the time. There are few suppliers who would consciously enter into a contract that caused them to fail at this rate. In practice, the SLA, when entered into by entities who understand the process, defines a level of performance at which the supplier will accept the risk of failure. This means that the actual delivered performance should be better than the SLA threshold by an amount that defines the risk. The gap between the two is a statistical value that depends on the population of services, the measurement time and the rate at which failures occur. There is no single relationship between these parameters and that is where the value of a model enters into the business equation.

SLA MISTAKES

The recent history of SLAs is, unfortunately, littered with examples that did not lead to the expected outcome for either party. Common mistakes include:

- Setting the target thresholds to the average of a series of actual measures
- Setting such complex SLAs that custom billing systems are needed to operate them
- Setting unnecessary targets that affect the design out of proportion to the performance requirement

- Including measures that have never been used before and which have little, if any, impact on the service delivered to the end users of the services

- Specifying measures in such fine detail, or over such a short period, that it is impossible to achieve statistical significance in the measures (this makes the risk almost unquantifiable)

The list is extensive, but all of the above apply to some degree in the assignments we have been involved with.

BALANCE OF RISK BETWEEN EQUIPMENT SERVICE PROVIDERS

DESIGN VS SERVICE CONSIDERATIONS

It is a common practice, particularly in major public procurements, that a service will be built and managed by different entities within a consortium of companies assembled specifically for the purpose through a Special Purpose Vehicle, or SPV. One member of the team will concentrate on the design, build and setting to work of the underlying system, while another concentrates on the ongoing service management of the new infrastructure. In public infrastructure procurements, the period of operation can be up to 25 years. This is very much longer than the period over which the equipment supplier needs to be involved. This leads to internal contractual pressures within the consortium.

The equipment designer has relatively short-term goals. As far as performance is concerned, they need to deliver the capacity and MTBF required by the design. This member wishes to handover the equipment at the earliest possible moment and then move on to other projects.

The service provider, on the other hand, is there for the long-term; often accepts ownership of the equipment from the supplier and is responsible for the maintenance and repair of the equipment. This member generally also accepts the risk of the SLAs and is highly vulnerable to the late emergence of a design defect long after the original supplier has no further interest in the project.

DESIGN SCRUTINY

The design process cannot be carried out in isolation by one of the parties, as Availability, to name but one parameter, is a combination of MTBF and MTTR that are delivered by different members of the consortium. Well-constructed consortia resolve this design conundrum amicably, while others, operating on a more rigidly contractual basis, have great difficulty in coming up with efficient solutions.

The methods of addressing this that have met with some degree of success include:

- Performing a major handover test that makes the equipment supplier prove that the SLA constraints have been met. This can be a very time-consuming process
- Withholding some part of the costs until the long term performance has been satisfactorily proven
- Imposing a long term duty of correction at their cost on the equipment supplier in the event that failures of design emerge at a later stage

The interest of all parties is to enter into a contract and accept responsibilities with their eyes open. The use of the model, in this context, is paramount. It can be of equal value to both sides in their negotiation. In this context, it is a means of achieving a contractual position that is equitable to both parties.

SENSITIVITY TO PERFORMANCE

MODEL BACKGROUND

It is possible to model the effects described in the preceding sections. The modelling discipline provides a highly effective mechanism for a systematic consideration of all aspects. The approach used by Intercai is to balance the detail and accuracy of the model with that of the data that will be entered into it.

Intercai employs analytic as opposed to simulation models as they tend to be effective for the purpose, better bounded and run much more quickly. They are also much easier to understand and any inherent assumptions made, in the interests of simplicity and speed of implementation, are made much clearer.

The real value of such models is the ability to adjust the assumptions made about the input data and see if there are any surprising effects on the penalties imposed. If there are then it is possible to model that area more closely, or even design an alternative that avoids the problem altogether. If not then the model has served its purpose.

In an ideal world the assumptions / penalty consequences process would take place before the system is designed so that the results could be included in the design. In many cases this does not happen and the SLAs are set in the contract such that the system has to be designed to meet them at an acceptable risk level without any opportunity for trade-offs.

MODEL EXAMPLE

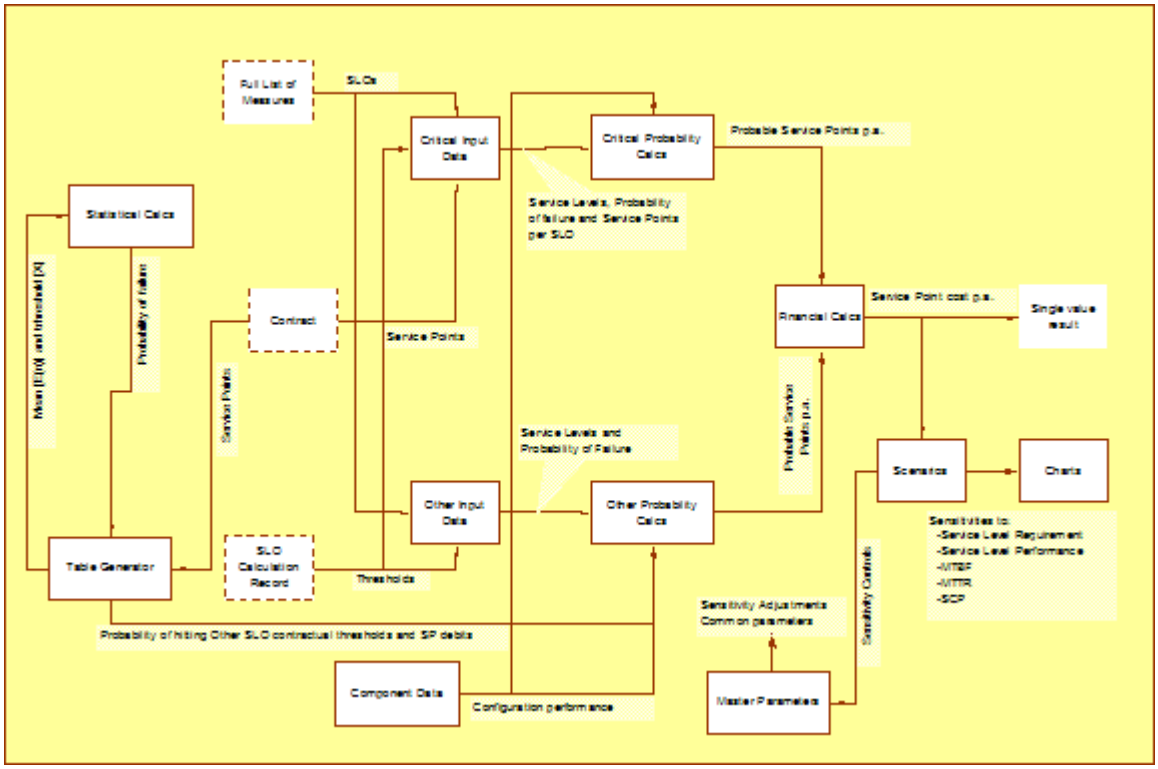

Figure 1 Typical Availability Model

An example from a recent Intercai model illustrates these points. Figure 1 outlines the data flows in a model that was built to support an IT outsourcing project.

In this case, the SLA was complex with something 100s of measures or Service Level Objectives (SLOs). The model allows for two separate penalty schemes, Groups 1 and 2, that are applied to different categories of SLO.

The model includes the expected performance of the different service configurations; the defined SLO thresholds for each service type and the definition of the penalty scheme. It works out the probability of failure for each service type and then calculates a weighted probability of failure and converts it into a monetary penalty value. This is the expected average penalty likely to be paid over the long term. The mathematics uses standard Poisson distributions as the means of modelling the random occurrence of faults, and is based on the Excel spreadsheet package.

MTBF SENSITIVITY

Once the model is complete and loaded with data it is possible to adjust the assumptions to see the effect. To make this a straightforward process, all of the input data includes a sensitivity parameter so that simply adjusting the sensitivity parameter adjusts all of the parameters of a single type at the same time. The individual parameters can be adjusted should this be necessary.

Some of the results are surprising. *Figure 2* shows the impact of moving the supplier's MTBF performance through a range of values. The nominal working point for the model is half way along the horizontal axis at the 100% point.

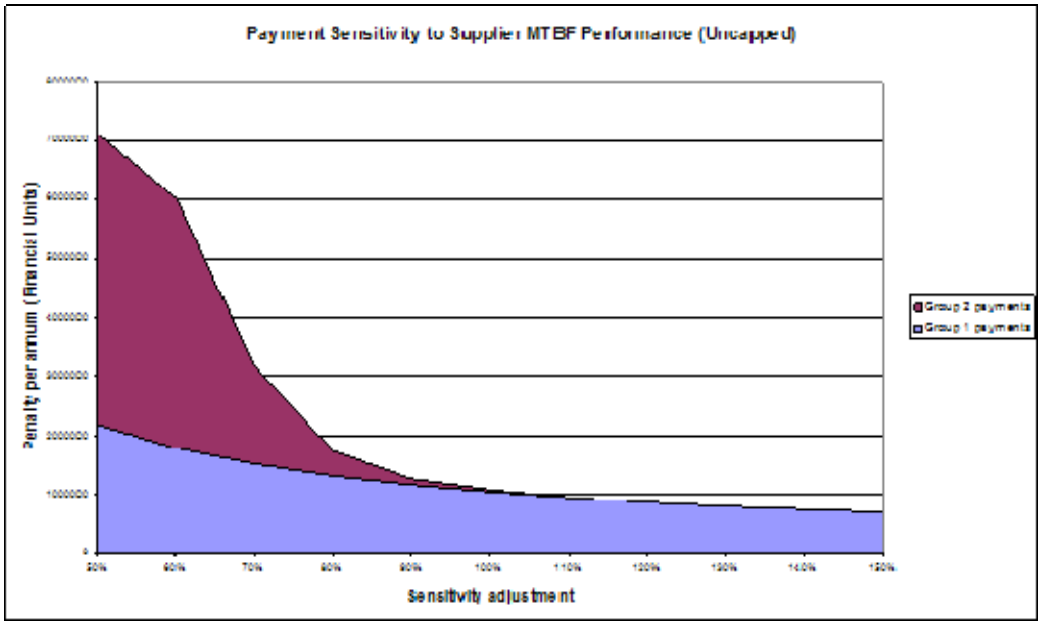

Figure 2 Example of Sensitivity

It can be seen that at the normal working point the payments amount to just over 1M financial units per annum with virtually all of that coming from the Group 1 services, less than 10% of the

total. As the MTBF is reduced, i.e. more faults per unit time, the Group 1 penalties increase slowly, but the Group 2, the 'non-critical ones', increase extremely rapidly to the point where they dominate. Indeed in this system, they took the payments a long way past the point where penalties were capped and so lost all incentive to improve the performance. Over a six-year contract this would have a major impact on the earnings of the project. It is interesting to note that at the working point the SLA system encourages the service operator to concentrate on the Group 1 SLOs. However, if the working point were in a slightly different place, the Group 2 penalties would dominate and produce precisely the opposite behaviour in the service operator than was intended.

AVAILABILITY SENSITIVITY

The same system exhibits an entirely different performance under testing the sensitivity against a different parameter. Figure 3 shows how the penalties vary with changes in the Availability target. The horizontal axis in this case is effectively exponential and it can be seen that for much of the working range there is no impact on the penalty at all.

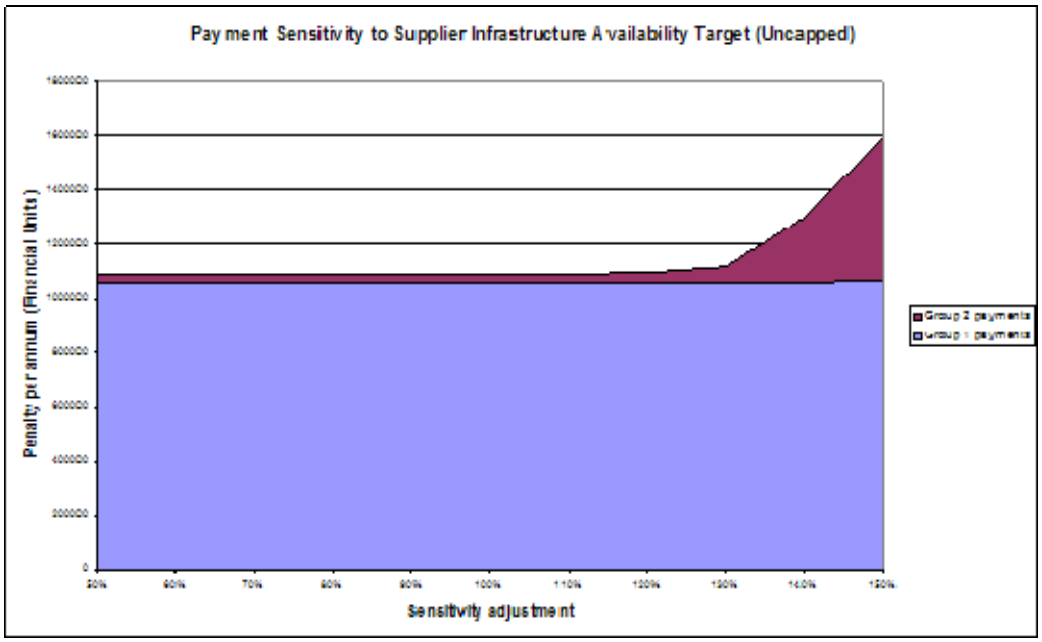

Figure 3 Example of insensitivity

There is a good reason for this and that is that the measurements were taken over a very short period relative to the rate of failure, which is once a month for each service individually. When a fault did occur, the repair time was well in excess of the average minimum allowable downtime, although all the intervening months had 100% Availability. This meant that the service operator could not affect the penalty payments in any way with regard to Availability adjustments. This is again an unintended consequence because the rational response to such an SLA is to make the repair strategy as cheap as possible, even at the expense of extended repair times. This is the result of failing to achieve statistical significance.

CONCLUSIONS

The modelling approach of Intercai provides a valuable insight into a number of key aspects of the provision and receipt of service. It enables the quantification of an SLA in an eyes open environment for both the supplier and the client. In the real world of outsourcing, the model serves a number of purposes:

- Providing management with an in depth understanding of the SLA regime for the whole outsourced system
- Enabling an understanding of the implications for the business that has outsourced its ICT
- Highlighting sensitivities, that is the SLA related aspects that have most effect
- Quantifying risk and exposure on both sides - user and supplier
- Assisting negotiation and in particular the use of a sensible number of SLAs that provide the right incentives for effective service provision
- Minimising financial exposure to the business
- Applicable to user and supplier sides regarding a number of roles including sales, commercial, operational and supplier